AAS 98-332

/N-33

INITIAL FLIGHT RESULTS OF THE TRMM KALMAN FILTER

Stephen F. Andrews*
Wendy M. Morgenstern*

The Tropical Rainfall Measuring Mission (TRMM) spacecraft is a nadir pointing spacecraft that nominally controls attitude based on the Earth Sensor Assembly (ESA) output. After a potential single point failure in the ESA was identified, the contingency attitude determination method chosen to backup the ESA-based system was a sixth-order extended Kalman filter that uses magnetometer and digital sun sensor measurements. A brief description of the TRMM Kalman filter will be given, including some implementation issues and algorithm heritage. Operational aspects of the Kalman filter and some failure detection and correction will be described. The Kalman filter was tested in a sun pointing attitude and in a nadir pointing attitude during the in-orbit checkout period, and results from those tests will be presented. This paper will describe some lessons learned from the experience of the TRMM team.

INTRODUCTION

TRMM Spacecraft

The Tropical Rainfall Measuring Mission (TRMM) spacecraft, seen in Figure 1, is a joint NASA/NASDA mission that was launched on November 27, 1997 from Tanegashima Space Center, Japan. The spacecraft is three-axis stabilized, in a near circular 350 km orbit at a 35 degree inclination. The Mission Mode is nadir pointing, and due to Sun constraints, the spacecraft must be rotated 180 degree about nadir (yaw) every few weeks. The sensor complement includes a static Earth Sensor Assembly (ESA), two two-axis Digital Sun Sensors (DSS), a redundant three-axis Inertial Rate Unit (IRU), eight Coarse Sun Sensors (CSS), and two Three-Axis Magnetometers (TAM). The spacecraft is controlled with four Reaction Wheels (RW), twelve thrusters (Reaction Engine Modules, REM), and momentum is unloaded with three Magnetic Torquer Bars (MTB). In Mission Mode, which is the nominal science configuration, attitude determination is done with the ESA for roll and pitch, and integrated IRU rate for yaw.

Problem Description

A potential single point failure of the ESA was first identified at Goddard Space Flight Center (GSFC) in 1992, with the discovery of a "fogging" effect of the ESA lenses[1]. This problem could cause the ESA to fail the Mission Mode attitude determination requirement. A backup attitude determination method was needed to satisfy the system redundancy requirements. Buying another ESA or a star tracker (ST) was not a realistic option, given the TRMM budget and schedule. A software backup using the available sensor measurements added redundancy without requiring additional hardware or affecting other subsystems such as power or structures.

---

*Aerospace Engineer, Guidance, Navigation, and Control Center, NASA's GSFC

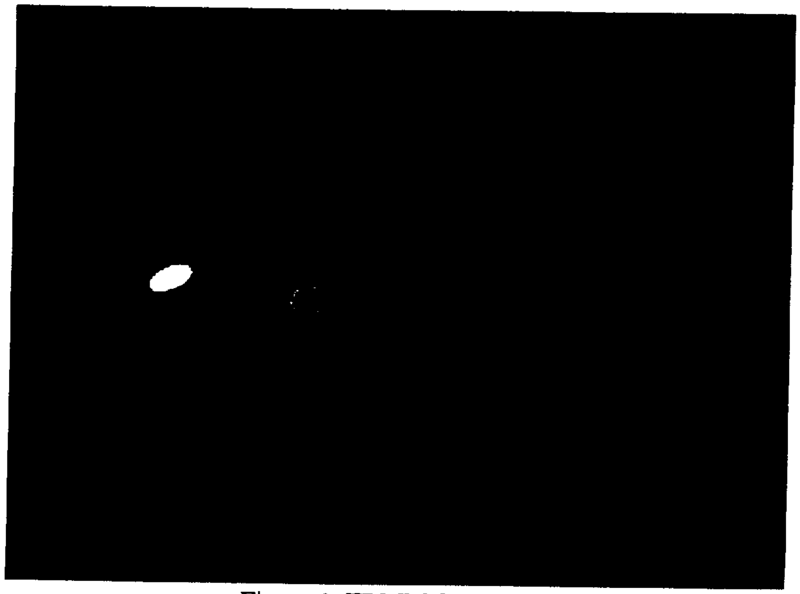

Figure 1 TRMM Spacecraft

ALGORITHM HERITAGE

A six state extended Kalman filter was chosen to backup the ESA based on the results of a trade study of several attitude determination algorithms[2]. The TRMM Kalman filter was adapted from the Rossi X-ray Timing Explorer (RXTE) Kalman filter, which was based on the Kalman filter developed for the Multimission Modular Spacecraft (MMS) by Murrell[3], and described in Lefferts, Markley, and Shuster[4]. The major changes to the RXTE algorithm were the replacement of the ST processing with TAM processing, the addition of a second DSS, and the coding of new initialization functions, associated subroutines, and new interfaces to fit the algorithm into the existing flight code[5]. The mathematical checks in the RXTE code were reused for TRMM, with modified limits. Since the core part of the Kalman filter code was already tested and working in flight on RXTE, there was a high degree of confidence in it which allowed a 'black box' testing approach to the new algorithm.

Typical spacecraft applications of a Kalman filter, such as on the Extreme Ultraviolet Explorer (EUVE) and RXTE, use a DSS and a ST to update the IRU propagated attitude. A DSS and a TAM were used in a TRIAD algorithm[6] for a deterministic attitude on the SAMPEX satellite. The TRMM Kalman filter updates the IRU propagated attitude and the IRU drift rate biases with TAM and DSS measurements. It is the first onboard attitude estimation algorithm at GSFC to use a DSS and a TAM to estimate spacecraft attitude to better than 0.25°.

ALGORITHM

The main portion of the TRMM algorithm is a discrete, extended Kalman filter.[7] For brevity's sake, only a few relevant points will be mentioned here.

The term residual as used in this paper refers to the difference between the vector measured by the sensor and the vector predicted by the model. The scalar implementation used in the flight software introduces another term that is called the adjusted residual. As each measurement component is processed, the vector predicted by the model is updated.

As a result, subsequent residual components have to be corrected for that update. These corrected residuals are referred to as adjusted residuals. The corrections are usually so small that the values of the residual and the adjusted residual are almost identical.

The algorithm also includes checks on the data in the filter. The first check is made on the availability and quality of the sensor data. For example, if the sun is in the DSS field of view but the measurement is not valid, the filter will not use that DSS measurement. In addition, there is a residual tolerance test that rejects any measurements that create residuals larger than a set tolerance. These checks prevent the estimation from using bad sensor data. This is not an algorithm failure, so no corrective action is taken.

There are three Failure Detection and Correction (FDC) checks designed specifically to monitor the Kalman filter algorithm. Two checks monitor the covariance matrix for divergence and positive semidefiniteness. The third test ensures that the adjusted residual remains within 3σ of the expected value of the residual. For all three tests, the ACS software autonomously performs the same actions. First, the software stops updating the attitude quaternion and the gyro drift with the failed sensor, and then it commands the spacecraft to a power and thermal safe attitude after a specified amount of time.

GROUND TESTING

Using software to add redundancy is not a trivial task. For TRMM, the most difficult issue was adding the new Kalman filter code into tested flight software without altering the existing ESA-based controllers. In addition, there were concerns about processor speed. Since it was unclear if the hardware could run fast enough to simultaneously process the ESA information and run the Kalman filter algorithm, the attitude control system (ACS) team decided to run one algorithm at a time.

A more detailed description of the software implementation and testing can be found in Andrews and D'Agostino[5]. Tests were run to verify the nominal performance of the Kalman filter, and to ensure that the existing control modes were not affected by the addition of the new algorithm. All test results were nominal, except one. In that test, the filter rejected DSS measurements after an eclipse. The DSS residuals passed the initial tolerance test, but failed the adjusted residual test, causing the filter to reject the DSS data. This indicates that either the DSS tolerance was set too tight or that the covariance did not grow large enough during eclipse, leaving the filter knowledge of the TAM noise smaller than it should have been. At the time, the ACS team believed the failure was the result of a mismatch between the 'true' ephemeris and the 'modeled' ephemeris in the test setup.

FLIGHT RESULTS

Sun Acquisition Mode Test

On the second day of the mission, TRMM was still in Sun Acquisition Mode holding the spacecraft x-axis 16.5 degrees from the sunline. In this mode, the spacecraft is controlled directly off the CSSs and the IRUs; the Kalman filter output is not used in the control loop. The Kalman filter was run for a total of 13000 seconds. After converging for 9640 seconds, the filter was reinitialized during eclipse to study the TAM-only filter performance. The TAM residuals for the entire test are shown in Figure 2. The flat line portions in this figure are periods of loss of signal (LOS), when TRMM was not in contact with the ground. It is obvious that the TAM residuals are not the zero mean, white noise processes modeled by the filter equations. The magnitude of this modeling error has

implications for setting the proper sensor noise parameters in the filter that will be discussed later.

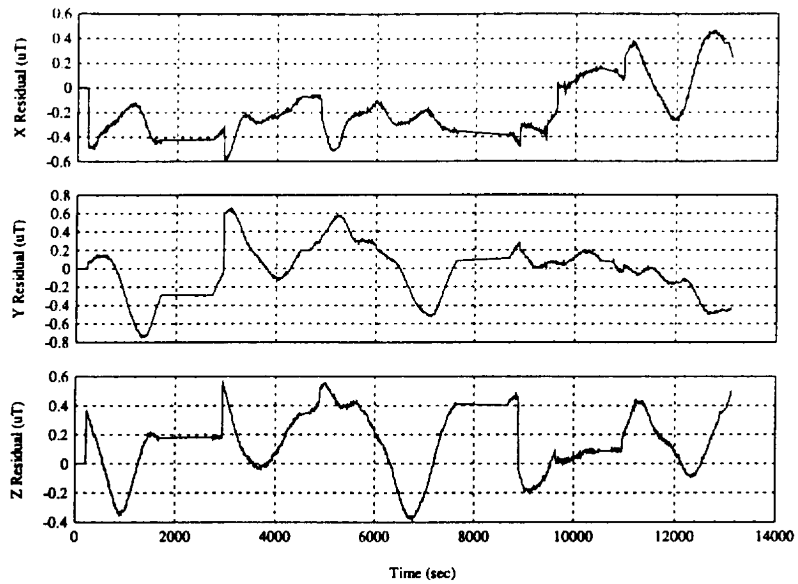

Figure 2 Sun Acquisition Mode Test: TAM Residuals

The standard deviations of the attitude estimate, shown in Figure 3, converged to [0.02,0.006,0.002] degrees within 8000 seconds. In Sun Acquisition Mode, the sun is held in the same location in the body frame, perpendicular to the z axis and primarily along the x axis. This reduces observability in the x axis, a phenomenon that is reflected in the relative size of the attitude standard deviations.

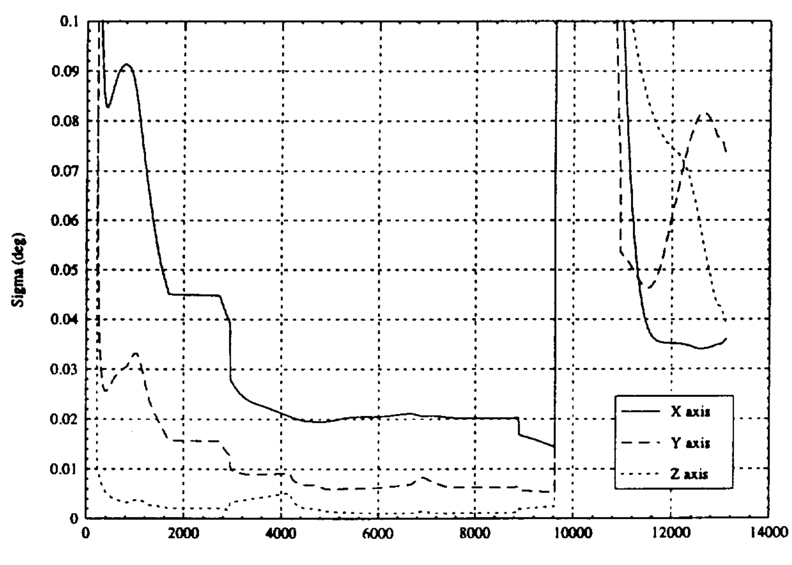

Figure 3 Sun Acquisition Mode Test: Standard Deviation of the Attitude Estimate

The standard deviations of the gyro bias estimates are shown in Figure 4. As with the attitude standard deviations, the relative sizes of the bias standard deviations indicate the measurement geometry. The value of the gyro bias estimate shown in Figure 5 reached steady state before the filter was reinitialized. The impact of the reset can be seen clearly, and the bias estimate was still reconverging when the test ended.

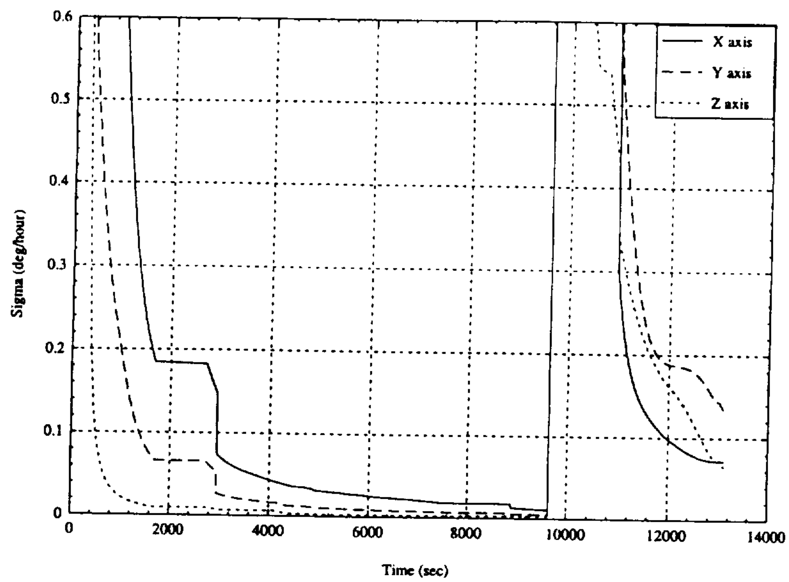

Figure 4 Sun Acquisition Mode Test: Standard Deviation of the Gyro Bias Estimate

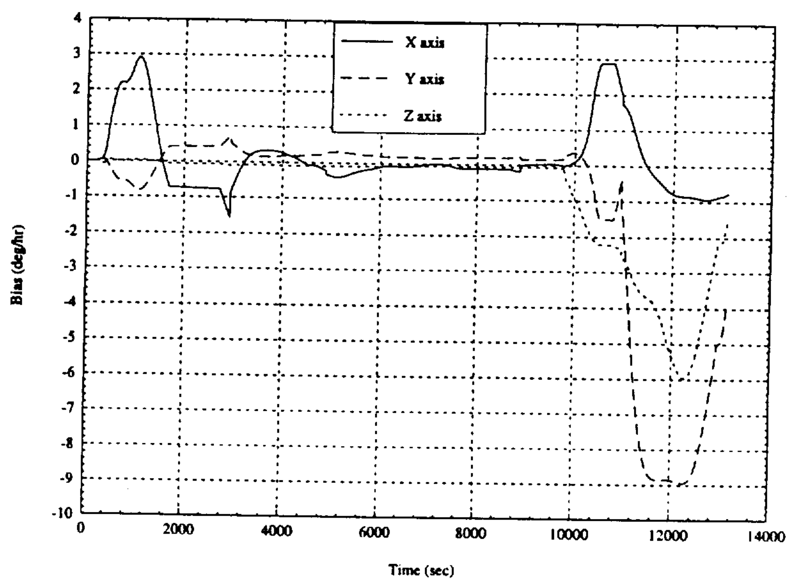

Figure 5 Sun Acquisition Mode Test: Estimated Biases

The DSS residuals are shown in Figure 6. When the Kalman filter is first enabled during sunlight, the DSS residuals were well within acceptable limits. When TRMM entered sunlight about 300 seconds after the filter was reinitialized, the DSS adjusted residual failed the 3σ tolerance check (~0.15 degrees) after less than 700 seconds in daylight. The adjusted residual continued to fail its tolerance check for the entire daylight portion of the orbit. TRMM went into eclipse before the DSS measurement was accepted again, and the test was ended.

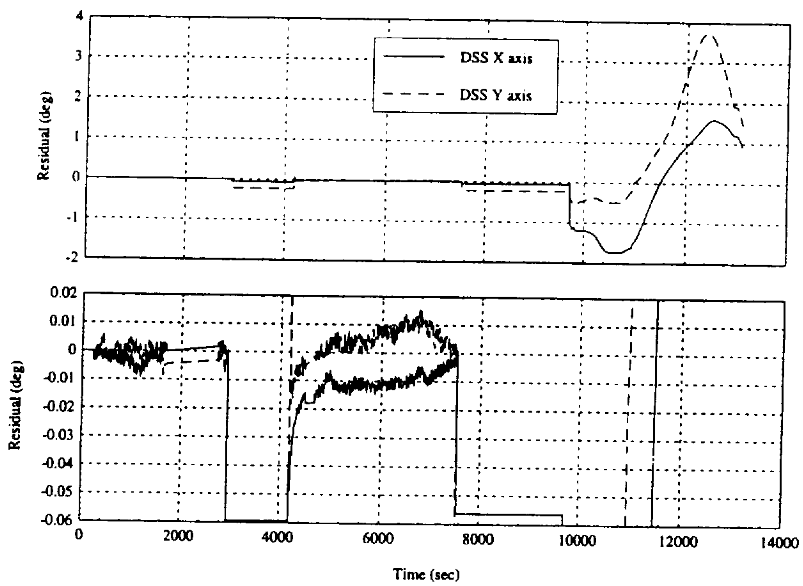

Figure 6 Sun Acquisition Mode Test: DSS1 Residuals

This adjusted residual test failed because the actual DSS adjusted residual was larger than the expected adjusted residual that is calculated from the 3σ tolerance, the state covariance, and the sensor noise. Review of the data showed that the state covariance was too small and the 3σ tolerance was too tight. The state covariance was too small because the Kalman filter was overweighing the TAM measurement, and converging too quickly. This weighting factor is a function of the TAM measurement noise covariance matrix that was set to model sensor noise on a zero mean process. The actual TAM measurement residual has a nonzero mean due to modeling errors. The 3σ tolerance was set too tight because the filter should be allowed to accept 5σ DSS data since the DSS was performing better than expected.

Other problems were identified later, after GSFC's Flight Dynamics Facility had time to analyze several days of flight data. It was found that the DSS heads were misaligned by as much as 0.3 degrees, which caused biased DSS residuals, leading to biased estimates. Also, the influence of the MTBs on the TAM measurements had not been accurately compensated for, and that increased the TAM residuals. In addition, the IRU calibration maneuvers had not been done yet, and the alignment matrices on-board did not properly account for the true IRU alignments. Finally, it was found that the magnetic field model on board was not internally consistent. The coefficients were from a 1995 model, but the epoch time for computation of the secular variations was set to 1990. This means that the residuals between the magnetic field model and the TAM measurements had a much larger bias and variance than expected.

Mission Mode Test

Before the next Kalman filter test, the filter parameters were retuned by modifying the on-board tables. First, the TAM measurement noise covariance matrix was increased from 0.25 μT^2 to 1.0 μT^2 to try to account for the model errors. Second, the DSS adjusted residual tolerance was increased from 3σ to 5σ. Third, the magnetic field model coefficients were set to the 1990 values to match the epoch time. Fourth, the DSS parameters were updated to account for some of the misalignment errors. The on-board software precluded compensating for the DSS misalignments completely, so there were still unmodeled DSS misalignments of up to 0.08 degrees. Once the changes were made, the Mission Mode test began.

The TAM residuals are shown in Figure 7, and, as in the Sun Acquisition Mode test, they are neither zero mean nor Gaussian distributed. The high frequency component of the signal is due to the unmodeled 0.5 Hz rotation of one of the payload instruments, and the low frequency variation may be due to the effects of the MTBs on the TAM measurements. The sharp spikes on the plot are caused by the on-board magnetic field model.

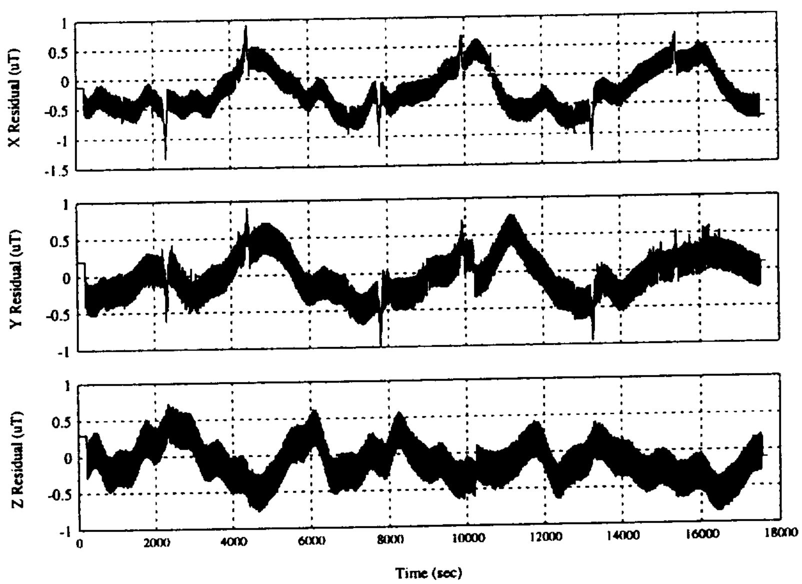

Figure 7 Mission Mode Test: TAM Residuals

The standard deviation of the attitude estimate is shown in Figure 8. The x/z (roll/yaw) quarter-orbit coupling is due to the one revolution per orbit rotation of the spacecraft about the y (pitch) axis. The spacecraft y axis is generally perpendicular to the sunline, and thus shows the greatest estimated accuracy. During eclipse, the covariance increases because the less accurate TAM is the only update sensor available. Figure 9 shows that the gyro bias estimate is also affected by the availability of the DSS measurement. The periods when the bias covariance is increasing or holding steady are periods of eclipse.

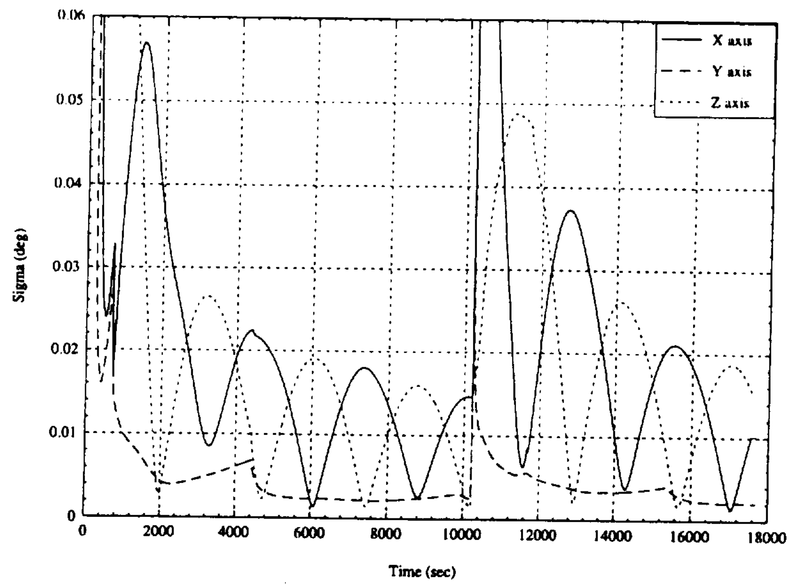

Figure 8 Mission Mode Test: Standard Deviation of the Attitude Estimate

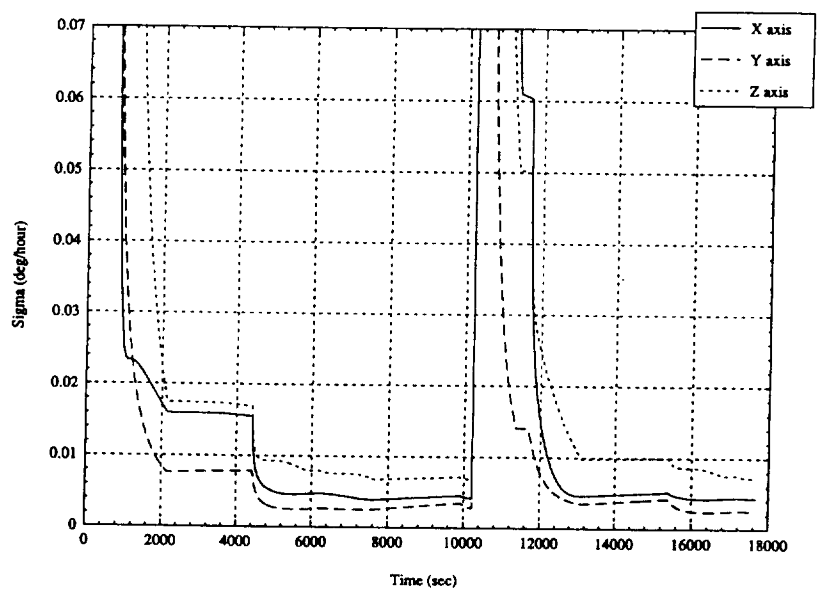

Figure 9 Mission Mode Test: Standard Deviation of the Gyro Bias Estimate

The Kalman filter estimated gyro biases are shown in Figure 10. Upon initialization and reinitialization, the initial attitude transient lasts about 2000 seconds. The filter bias estimate settles to the same values before and after the reinitialization; this shows that the gyro drift rate is steady on a time scale of hours. This result is expected because of the high quality and drift stability of the TRMM IRUs.

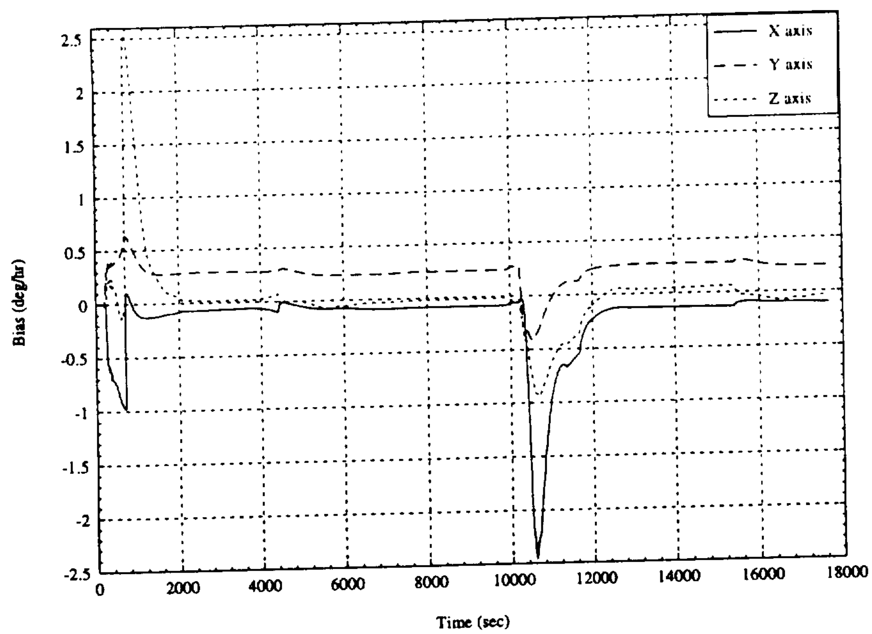

Figure 10 Mission Mode Test: Estimated Biases

The DSS residuals are plotted in Figures 11 and 12. Between measurements, the filter simply stores the last value of the residual, and the data goes static. The effect of the DSS misalignments can be seen in the large initial values of the residuals when the sun first enters the DSS field of view.

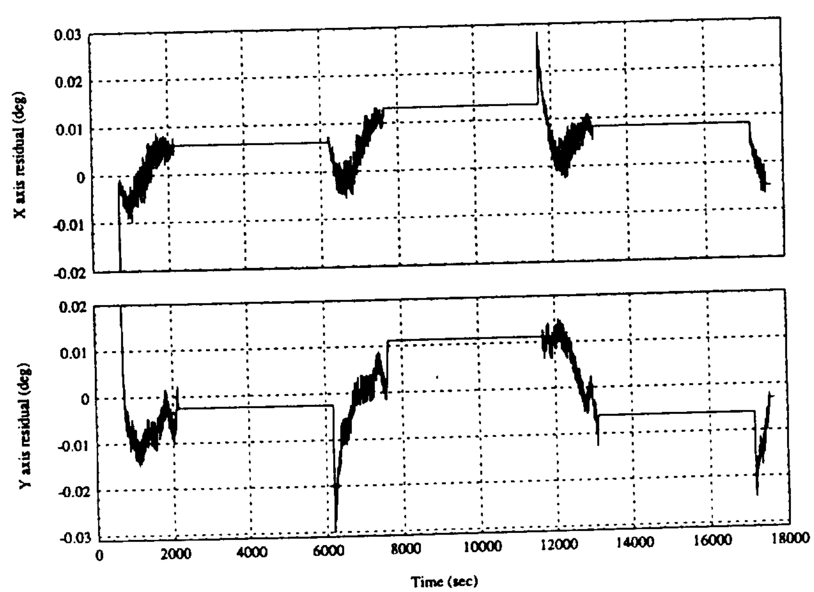

Figure 11 Mission Mode Test: DSS1 Residuals

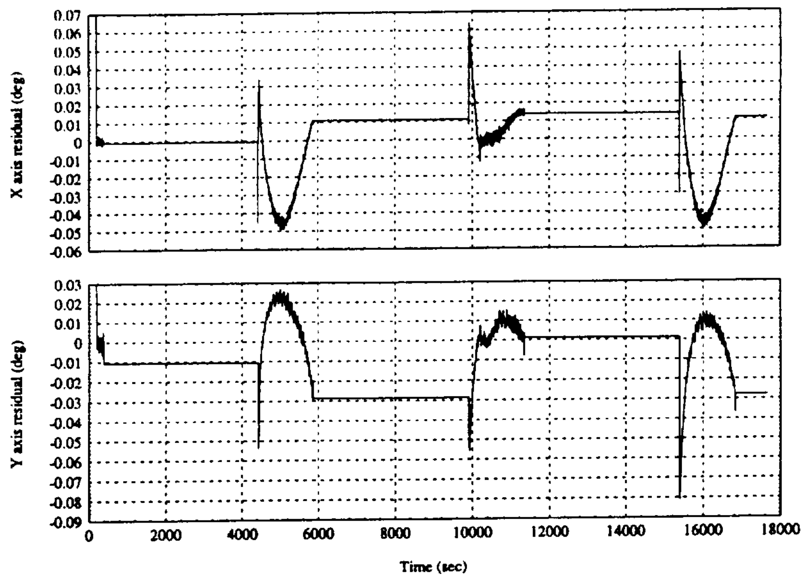

Figure 12 Mission Mode Test: DSS2 Residuals

Figure 13 Mission Mode Test: ESA Attitude

The best measure of the performance of the Kalman filter is the attitude derived from the ESA. Although the ESA data is not processed on board when the Kalman filter is running, the unprocessed data is available in telemetry. With this information, the ESA attitude was calculated on the ground and is shown in Figure 13. The initial attitude

transients are on the order of 1° in roll, and 0.2° in pitch. The best performance of the filter is the period from 5000 to 10000 seconds, when the largest attitude error is about 0.12°.

LESSONS LEARNED

The possible improvements in the TRMM Kalman filter fall into two basic categories: operational issues and performance issues. Operational issues include transitions to the backup mode, testing, software design, and data flow. Performance issues pertain mostly to properly tuning a Kalman filter so that it functions effectively with real sensors.

Operations issues

Since the TRMM Kalman filter was added late in the testing cycle, the concern about onboard processing power forced the software design to an either/or mindset. Either the ESA processing could be run or the Kalman filter could be run, but not both. The filter was designed under the assumptions that the ESA had failed and that once the Kalman filter was turned on, it would never be turned off. It was also assumed that the Kalman filter would have to replace only the ESA functions, such as the earth acquisition maneuvers, maintaining nadir pointing, and inertial slews and holds.

Assuming the filter only had to replace the ESA functions meant that the filter's performance during other cases, such as thruster maneuvers, was not thoroughly considered. This led to several oversights in the filter design. First, in all of the ACS control modes, except during thruster maneuvers, the ACS software runs at a 2 Hz cycle. During thruster maneuvers, the controller runs at 8 Hz. However, since the Kalman filter was only coded as a replacement for ESA functions, the 2 Hz duty cycle was hardcoded into the Kalman filter algorithm, and the filter cannot run during the thruster maneuvers. Second, the software propagates the attitude estimate during thruster maneuvers, but it does not propagate the filter covariance. Thus, after completing the thruster burn, the covariance gives an incorrect indication of the accuracy of the attitude estimate. A solution is to reinitialize the filter and allow it to reconverge, a process the flight tests show takes several hours, which reduces the quality of the science data for that period of time. Finally, since normal operations, such as the Delta-V maneuver, require the filter to be reinitialized periodically, testing all possible reset conditions should be included in both ground and flight testing, a luxury the TRMM schedule did not allow. More thorough testing might have revealed more of these problems in time for the development team to modify the design, rather than forcing the operations team to resort to work-arounds.

Since it was assumed the Kalman filter would never be turned off, the ground verification did not test the transition between the filter and ESA processing. Again, this led to several oversights. The first problem concerns the gyro biases. The filter is continually estimating the gyro biases for all three axes. When the filter stops running, these bias estimates are stored in memory. As the spacecraft's orientation changes during nadir pointing, misalignment errors map differently into the gyro drift bias error. If the filter is running, these changes will be compensated for on board. If the filter is not running, the estimate that is in memory may actually introduce a small error in the drift biases. Thus, it is necessary to reset the gyro drift biases after exiting the filter. However, the initial estimated bias is set in the gyro initialization subroutine, not in the Kalman filter initialization/reset subroutine. Commanding a filter reset only reruns the filter initialization. To zero the estimated bias, the gyro initialization subroutine must be rerun. This can only be accomplished by rebooting the ACS software, which is extremely risky to do in flight.

To eliminate this risk, a separate gyro reset command had to be incorporated into the ACS software a few months prior to launch. Including the zeroing of this initial biases in the filter initialization subroutine would have been a much cleaner solution.

These issues make the transition between ESA attitude determination and Kalman filter attitude determination unnecessarily awkward, and complicated the on-orbit testing. The availability of the Kalman filter data is another operational problem that resulted from adding a backup algorithm to a mature software design.

Currently, there are problems getting flight data from the Kalman filter because the telemetry packet is available only by special request, or asynchronously. This means that the operations team must send a new command to the spacecraft every nine hours to keep the Kalman filter data in the telemetry stream. In addition, the filter packet is only issued every eight samples, which is insufficient for a thorough performance evaluation if the filter ever becomes the primary attitude determination method. Since it is in an asynchronous packet, the flight recorder does not store the filter data, so real-time playbacks must be used to regenerate the data on the ground. This is inefficient, and requires a large effort from the Flight Operations team. The resulting data is full of gaps, since the Flight Operations team only records the telemetry stream during real time passes. Fortunately, the flight software allows the team to modify the data storage operations, so it is possible to record a continuous data stream for the filter information. Unfortunately, the data rate is still one sample in every eight, and it involves yet another operational workaround. Many of these data problems could have been avoided if the asynchronous packet had been redefined as a synchronous packet that is always available in the telemetry stream and is always sent to the flight recorder, and issued at a higher data rate.

Performance issues

The Kalman filter models assume zero-mean white noise measurement residuals, which is mostly true of the DSS residuals but is not true of the TAM residuals. The filter has no knowledge of biased sensor readings unless they are included in the state equations, so the DSS misalignment has a large impact on the accuracy of the filter. To characterize the estimation errors caused by instruments, the sensors and relevant instruments must be accurately modeled in the simulation including biases, scale factor errors, and misalignments. In particular, an accurate gyro model is essential if gyro biases are included in the filter states. Also, the full effects of the Earth's magnetic field on the Kalman filter cannot be properly seen in simulation because the low frequency variations of the Earth's magnetic field are hard to model accurately. Ideally, the simulation should model all of the errors that will be seen on orbit, but that is not always easy to achieve.

The on-orbit test needs to be run for many hours to adequately test the backup algorithm. Ideally, the filter should be tested under all the conditions where it is expected to be used. As with all on-orbit tests however, this requirement has to be balanced with other subsystem tests and the science schedule.

As mentioned previously, FDC is designed to capture certain problems, and to keep the spacecraft safe. The three FDC tests discussed will not indicate if the filter is trying to estimate an attitude error or a gyro bias error larger than it was told to expect. If the true error is larger than the error indicated by state covariance, the filter may converge to an incorrect attitude. This type of error is indicated when the filter's estimates do not match the true attitude. Outside of computer simulation, however, there is no truth model to use for comparison. If this type of problem is suspected, the best option is to compare the

estimates from the filter to ground estimates and adjust, or tune, the filter based on the comparison, a luxury that is not always available to the operations team. To properly tune a filter, the designers must be aware of the largest possible state error, and choose initial state covariance values larger than expected. Even though this greatly increases the filter's convergence time, it will allow the filter to converge from larger, more uncertain initial estimation errors. This is obviously a filter design tradeoff.

Implementation features

The flight software developers should rarely hardcode a number; a table design that allows parameters to be changed with a simple uplink rather than a software patch should be used instead. Software patches require a significant development and testing effort from the software maintenance team, and risk the safety of the spacecraft. For example, some of the DSS misalignment error was calibrated out of the data by changing some table values, but updating the magnetic field model to a 1995 epoch will require a software change.

One good feature of the flight software design is that it allowed the flight operations team to safely verify the filter's performance. There should be a control mode available to check out the Kalman filter performance before controlling with the filter's attitude estimates. On TRMM, the Sun Acquisition control mode uses the CSSs and IRUs for attitude determination, which allowed the filter to be tested in-flight without affecting the safety of the spacecraft. If the processing power had been sufficient to run both the Kalman filter and the ESA processing, much of the awkward testing done on TRMM would have been unnecessary. Running both algorithms simultaneously would provide two attitude estimates at all times, allowing ground personnel the luxury of evaluating the long term performance of the filter without affecting nominal mission operations.

A vital safety feature of the TRMM design is the FDC logic The three FDC tests pertaining to the Kalman filter allow the on-board algorithm to determine when it is inappropriate to use the filter results. The TRMM design stops updating the filter and autonomously places the spacecraft in a power safe mode if a bad attitude estimate is computed.

With approximately six months to go from a trade study to completely tested flight software, time constraints made it impossible to test the Kalman filter under every possible flight condition. Better system engineering should be able to identify possible failures and available backups early in the design phase. Flight results and trend data from each component should be reviewed early in the design to identify potential failures. Decisions to incorporate backups for hardware failures, whether using redundant hardware or implemented software backups, should be made early in the program. For TRMM it would have been best if the backup mode had been included in the earlier design, so that the additional processing needs would be reflected in the processor requirements and design. Once backups algorithms are selected, good subsystem engineering should help identify all possible uses of these algorithms so those conditions can be tested. 'Expected' usage tests do not cover all reasonable situations. The fundamental lesson learned is that software designed for one spacecraft can be reused on a different spacecraft if the software design is modular enough and the reused software is well tested, even though it is easy to underestimate the difficulty of the conversion.

CONCLUSION

In the unlikely event of a complete ESA failure, TRMM can meet pointing

requirements by using TAM and DSS data. The additional software in the form of an extended Kalman filter saved the expense of providing a star tracker for backup. It was not easy to convert existing software, but the software modifications were completed quickly, largely due to the modular design of the RXTE and TRMM software. Pre-flight testing of the software was extensive, but it did not find all the problems, both because there was not enough time to modify the simulation system to include all known error sources, and because some unexpected errors showed up in flight tests. However, flight testing showed pointing performance better than the required 0.7° and approaching the 0.2° performance of the primary attitude control system.

REFERENCES

1 D. K. Ward, internal memo, "TRMM ACS Peer Review Action Items Regarding Earth Sensor Assembly Concerns", August 18, 1992.

2 J. L. Crassidis, S. F. Andrews, F. L. Markley, K. Ha, "Contingency Designs for Attitude Determination of TRMM", Flight Mechanics/Estimation Theory Symposium, 1995 pp. 419-433.

3 J. W. Murrell, "Precision Attitude Determination for Multimission Spacecraft, "AIAA Paper 78-1248, Aug. 1978.

4 E. J. Lefferts, F. L. Markley, and M. D. Shuster, "Kalman Filtering for Spacecraft Attitude Estimation," *Journal of Guidance, Control and Dynamics*, Vol. 5, No. 5, Sept-Oct. 1982, pp 417-429.

5 S. F. Andrews, J. M. D'Agostino, "Development, Implementation, and Testing of the TRMM Kalman Filter," Flight Mechanics Symposium 1997, pp. 457-471.

6 J. S. Wertz, (ed.), *Spacecraft Attitude Determination and Control*, D. Reidel Publishing Co., Dordrecht, The Netherlands, 1984.

7 A. Gelb (ed), *Applied Optimal Estimation,* The M.I.T. Press, Cambridge, Massachusetts, 1974.